

Requirements for Evaluation of Monitoring System for CDF Offline

Gerald Guglielmo

(CD-doc-2653)

Abstract

As part of the initial investigations for the CDF Initiative to improve operational support load of the CDF Offline CAF systems, it was recognized that deployment of a monitoring system for the CAF would be vital early on in the process. To evaluate and select an appropriate monitoring system a set of requirements need to be defined. This document defines the set of requirements to be used in the evaluation stage.

Introduction

This document describes the requirements to use for evaluation of a monitoring system for CDF Offline as part of the CDF Initiative. The requirements put forth in this document are intended for the selection of a monitoring system that can perform the tasks needed for improving the current monitoring of CAF operations at CDF. These requirements are limited in scope to this specific task and not intended as a general set of requirements for monitoring systems, however by nature they could be sufficient for other experiments or similar operational environments.

Evaluation Requirements

The requirements as presented here are divided into two categories: absolutely necessary; negotiable if not practical. As defined the first category is composed of must have features which cannot be compromised on as the system will simply not perform the tasks needed without them. The second category are essential requirements for the first pass through on evaluation, but if no system was found to meet them those requirements could be relaxed.

Absolutely Necessary

- Web based interface for monitoring – The ability to monitor the system from the web is vital in the environment where operations support personnel are not manning a control room 24 hours a day, 7 days a week where dedicated consoles could be used.
- Ability to easily graph data – Monitoring information provided in integer or float format should be available for graphing by the monitoring system. Restricting the ability to define graphs to an individual, or set of experts, is acceptable and may be desirable for performance and reliability reasons.
- Ability to store and 10's of Gigabytes of monitoring data – The infrastructure for storing and later allowing access to monitoring data must be able to support at least 10's of Gigabytes of data.
- Ability to configure data retention period – The infrastructure must allow for configuration of the data retention period so that this can be tuned to the needs providing a high level of support while minimizing the support load on personnel. There is no set period of time for data retention that meets the needs of all situations and all goals.
- Capability of storing monitoring information for thousands of hosts – If the CAF Operations team decides to monitor worker nodes, then currently the number of nodes that would need monitoring is about 500. It is possible that in the next few years this will expand, but even if it does not it is best to not be near the limit of what the system can

support. Thus the ability to monitor thousands of nodes provides some margin of safety.

- Ability to track 10,000 different items – It is estimated that if monitoring of worker nodes is done there will be a need for several thousand items to be checked overall. Therefore a requirement of 10,000 means the system should be able to perform as needed without being stressed to the limit.
- Ability to define the frequency of data checks – There is no one frequency that works for all situations when it comes to monitoring systems. In some cases one might need near real time information, and in others having occasional checks is sufficient. There is no need for real time checking, but the ability to vary the frequency as needed is important. At the very least one should be able to vary the frequency from once per minute to once every 10 minutes.
- Ability to define sets of tests that can be applied to hosts – Not all hosts provide unique services that require specialized monitoring. In fact, in many cases the opposite is true and there is great benefit in being able to define sets of tests one time and have them applied to multiple hosts.
- Ability to group different types of hosts into a logical unit – While some hosts may provide similar services or need similar checks, they may not logically belong to the same system that provides a service. Many times systems are composed of many hosts that each provide different services, and thus require different monitoring, but should be grouped together for assessing the overall health of the system as a whole.

Very Important But Negotiable

- Web based monitoring by general users around the world – This would provide a boost in the ability to catch problems early on before they progress to a crisis stage by leveraging the willingness of users of CAF resources to check on the systems. While the potential gain here is large, if it became necessary to restrict some aspects of monitoring to a subset of users this would not be catastrophic.
- The ability to accept input data from agents in text form to standard output – Existing infrastructure that provides valuable monitoring information may already exist. Often they provide information on standard output, for example many of the Linux command line interface commands, and these could immediately be integrated into a new system without change if processing of standard output was supported.
- Ability to leverage existing expertise at the lab – Installing and supporting unique systems requires extra effort to understand and resolve issues in the underlying system. The ability to tap into an existing knowledge base or leverage effort experienced personnel can be a great asset.